

120
AF

Docket No.: 42P13499

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

In re Application of: Hubbert Smith

Assignee: Intel Corporation

Examiner: Ho, Thang H.

Application No.: 10/080,063

Art Group: 2188

Filed: February 19, 2002

For: Network Data Storage-Related
Operations

RECEIVED

SEP 27 2004

Technology Center 2100

APPEAL BRIEF
IN SUPPORT OF APPELLANT'S APPEAL
TO THE BOARD OF PATENT APPEALS AND INTERFERENCES

Mail Stop: Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Appellant hereby submits this Appeal Brief in support of Appellant's Appeal from final rejection of the pending claims in the above-captioned case.

A Notice of Appeal was timely filed on July 16, 2004.

The fee for submission of this Appeal Brief accompanies this Appeal Brief. A Petition For a One Month Extension of Time for submission of this Appeal Brief also accompanies this Appeal Brief, together with the fee required for said Petition.

An oral hearing is NOT desired.

Appellant respectfully requests consideration of this Appeal by the Honorable Board of Patent Appeals and Interferences, and allowance of the claims of the subject application.

Please charge any fees and/or credit any overcharges to Deposit Account No. 50-0221.

10080063
09/24/2004 SSESHE1 00000025 500221
01 FC:1402 330.00 DA

I. REAL PARTY IN INTEREST

The subject application is assigned to, and the real party in interest is, Intel Corporation of 2200 Mission College Boulevard, Santa Clara, California 95052.

II. RELATED APPEALS AND INTERFERENCES

To the best of the knowledge of the Appellant, the undersigned attorney, and the Assignee of the subject application, there are no prior or pending appeals, interferences, or judicial proceedings which may be related to, directly affect or be directly affected by, or have a bearing on the Honorable Board's decision in the subject appeal.

III. STATUS OF THE CLAIMS

Claims 1 - 43 were finally rejected in the Final Office Action mailed April 19, 2004, and are on appeal. No claims have been cancelled from the subject application. The Examiner confirmed the final rejection of claims 1 - 43 in an Advisory Action mailed June 10, 2004.

A copy of all of the claims on appeal is attached hereto as Appendix A.

IV. STATUS OF AMENDMENTS

In response to the Final Office Action, Appellant mailed on May 25, 2004 an Amendment After Final Rejection Under Rule 116 (hereinafter, "Amendment"). In the Advisory Action, the Examiner indicated that the Amendment would be entered by the Examiner. In response to the Advisory Action, Appellant timely filed a Notice of Appeal on July 16, 2004.

V. SUMMARY OF CLAIMED SUBJECT MATTER

Summary of the Independent Claims on Appeal

Of the claims on appeal, claims 1, 5, 10, 14, 16, 20, 25, 29, 31, 33, 37, 39, and 42 are independent claims. Independent claim 1 claims a first adapter for use in a first server in a network. In essence, the first adapter comprises a host bus adapter (HBA) that includes circuitry to cause, in response to a first request received by the HBA, the execution of a first data storage-related operation associated with a first set of mass storage devices. Additionally, the circuitry is capable of issuing, in response to the first request, a second request from the HBA to a second adapter in a second server in the network. The second request is to cause the second adapter to perform, in response to the second request, a second data storage-related operation associated with a second set of mass storage devices. (Claim 1; Specification, page 4, line 24 - page 20, line 26; Figures 1 - 9).

Independent claim 5 is directed to a first adapter to use in a first server. The first adapter comprises an HBA that includes circuitry to perform, in response to a request, a data storage-related operation associated with a first set of mass storage devices. The request is issued from a second adapter in a second server in response to another request received by the second adapter. The request is to cause the second adapter to perform, in response to the other request, another data storage-related operation associated with a second set of mass storage devices. (Claim 5; Specification, page 4, line 24 - page 20, line 26; Figures 1 - 9).

Independent claim 10 claims a first I/O processor. The first I/O processor is comprised in an HBA in a first server. The first I/O processor is configured to execute a set of operations. The set of operations includes the execution, in response to a first request, of a first data storage-related operation associated with a first set of mass storage devices. The set of operations also includes the issuance, in response to the first request, of a second request from the first server to a second I/O processor in a second server. The second request is to cause the second I/O processor to perform, in response to the second request, a second data storage-related operation

associated with a second set of mass storage devices. (Claim 10; Specification, page 4, line 24 - page 20, line 26; Figures 1 - 9)

Independent claim 14 is directed to a first I/O processor. The first I/O processor is included in an HBA in a first server. The first I/O processor is configured to execute a set of operations. The set of operations includes the execution, in response to a request, of a first data storage-related operation associated with a first set of mass storage devices. The set of operations also includes the issuance, after the completion of the execution of the first data storage-related operation, of a first message from the first server to a second I/O processor in a second server. The first message is to cause the second I/O processor to issue, in response to the first message, a second message to a process in the second server. The second I/O processor is operatively configurable both to generate the request and to perform a second data storage-related operation associated with a second set of mass storage devices, in response to another request from the process. (Claim 14; Specification, page 4, line 24 - page 20, line 26; Figures 1 - 9)

Independent claim 16 claims a method of using a first adapter. The first adapter includes an HBA in a first server in a network. The method comprises using the first adapter to perform, in response to a first request, a first data storage-related operation associated with a first set of mass storage devices. The method also includes issuing from the first adapter, in response to the first request, a second request to a second adapter in a second server in the network to cause the second I/O adapter to perform, in response to the second request, a second data storage-related operation associated with a second set of mass storage devices. (Claim 16; Specification, page 4, line 24 - page 20, line 26; Figures 1 - 9)

Independent claim 20 recites a method of using a first adapter. The first adapter includes an HBA in a first server. The method comprising using the first adapter to perform, in response to a request, a data storage-related operation associated with a first set of mass storage

devices. The request is issued from a second adapter in a second server in response to another request received by the second adapter. The second adapter is configured to perform, in response to the other request, another data storage-related operation associated with a second set of mass storage devices. (Claim 20; Specification, page 4, line 24 - page 20, line 26; Figures 1 - 9)

Independent claim 25 claims a method of using a first I/O processor. The I/O processor is in an HBA in a first server. The method comprises using the first I/O processor to execute, in response to a first request, a first data storage-related operation associated with a first set of mass storage devices. The method also includes using the first I/O processor to generate, in response to the first request, a second request. The second request may be issued from the first server to a second I/O processor in a second server. The second request is to cause the second I/O processor to perform, in response to the second request, a second data storage-related operation associated with a second set of mass storage devices. (Claim 25; Specification, page 4, line 24 - page 20, line 26; Figures 1 - 9)

Independent claim 29 claims a method of using a first I/O processor. The I/O processor is in an HBA in a first server. The method includes using the first I/O processor to execute, in response to a request, a first data storage-related operation associated with a first set of mass storage devices. The method also includes using the first I/O processor to generate, after the completion of the first data storage-related operation, a first message. The first message may be issued from the first server to a second I/O processor in a second server. The first message may cause the second I/O processor to issue, in response to the first message, a second message to a process in the second server. The second I/O processor may be operatively configurable both to generate the request and to perform a second data storage-related operation associated with a second set of mass storage devices, in response to another request from the process. (Claim 29; Specification, page 4, line 24 - page 20, line 26; Figures 1 - 9)

Independent claim 31 claims a network. The network includes a first server associated with a first set of mass storage devices. The first server includes an HBA that comprises a first I/O processor. The network also includes a second server remote from the first server. The second server is associated with a second set of mass storage devices. The second server includes a second I/O processor. Additionally, the network includes a network communication link that couples the first server to the second server. The first I/O processor is configured to cause a set of operations. The set of operations includes the execution, in response to a first request, of a first data storage-related operation associated with the first set of mass storage devices. The set of operations also includes the issuance, in response to the first request, of a second request from the first server to the second server via the link. The second request may cause the second I/O processor to perform, in response to the second request, a second data storage-related operation associated with the second set of mass storage devices. (Claim 31; Specification, page 4, line 24 - page 20, line 26; Figures 1 - 9)

Independent claim 33 claims computer program instructions residing in a computer-readable memory. The instructions include a set of instructions that when executed by a first processor result in a set of operations. The first processor is comprised in an HBA in a first server. The set of operations include the execution of, in response to a first request, a first data storage-related operation associated with a first set of mass storage devices. The set of operations also includes the generation of, in response to the first request, a second request that may be issued from the HBA to a second processor in a second server to cause the second processor to perform, in response to the second request, a second data storage-related operation associated with a second set of mass storage devices. (Claim 33; Specification, page 4, line 24 - page 20, line 26; Figures 1 - 9)

Independent claim 37 is directed to computer program instructions residing in a computer-readable memory. The computer program instructions include a set of instructions that when executed by a first processor result in a set of operations. The first processor is included in an HBA in a first server. The operations include the execution, in response to a request, of a first data storage-related operation associated with a first set of mass storage devices. The operations also include the generation, after completion of the execution of the first data storage-related operation, of a first message that may be issued from the first server to a second processor in a second server. When issued, the first message may cause the second processor to issue, in response to the first message, a second message to a process in the second server. The second processor is operatively configurable both to generate the request and to perform a second data storage-related operation associated with a second set of mass storage devices in response to another request from the process. (Claim 37; Specification, page 4, line 24 - page 20, line 26; Figures 1 - 9)

Independent claim 39 is directed to a first server. The first server includes an HBA that includes a first processor. The first processor is configured to be able to cause a set of operations. The operations include the execution, in response to a first request, of a first data storage-related operation associated with a first set of storage devices. The first set of storage devices is associated with the first server. The operations also include the issuance, in response to the first request, of a second request from the first server to a second server. When issued, the second request may cause a second processor in the second server to perform, in response to the second request, a second data storage-related operation associated with a second set of storage devices. The second set of storage devices is associated with the second server. (Claim 39; Specification, page 4, line 24 - page 20, line 26; Figures 1 - 9)

Independent claim 42 claims a first server. The first server includes an HBA that includes a first processor. The first processor is configured to be able to cause a set of operations. The operations include the execution, in response to a request, of a first data storage-related operation associated with a first set of storage devices. The operations also include the issuance, after completion of the execution of the first data storage-related operation, of a first message from the first server to a second processor in a second server. When issued, the first message may cause the second processor to issue, in response to the first message, a second message to a process in the second server. The second processor is operatively configurable both to generate the request and to perform a second data storage-related operation associated with a second set of storage devices, in response to another request from the process. (Claim 42; Specification, page 4, line 24 - page 20, line 26; Figures 1 - 9).

Conventional Techniques And Their Disadvantages

In one type of conventional computer network, one or more server computer nodes or “servers” are provided that may be associated with (i.e., exchange data with and control) one or more arrays of disk mass storage devices in which user data generated by processes executing in the network may be stored, and from which such data may be retrieved by such processes. At least one of these arrays may be a redundant array of independent disks (“redundant array”) in which a copy (or “mirror”) of a primary data volume stored in a primary array of disk mass storage devices (“primary array”) may be replicated, and from which the replicated data volume (“the redundant data volume”) may be recovered, in the event of a failure of the primary array. A redundant array is said to be “associated with” a primary array, if the redundant array stores a redundant data volume that is a copy of a primary data volume in the primary array. Also, a

redundant data volume is said to be associated with a primary data volume if the redundant data volume is a mirror of the primary data volume. (Specification, page 1, lines 3 - 16).

In this conventional network, a server (termed an “originating” server) that is associated with a primary array may be coupled via a network communication link to another server (termed a “target” server) that is associated with a redundant array, and the redundant array may be associated with the primary array. The originating server and the target server may each comprise a respective network interface card (NIC), HBA, central processing unit (CPU), and system memory. In each of the originating and target servers, the respective CPU and respective system memory are interconnected by a respective system bus, the respective NIC and respective HBA are interconnected by a respective I/O bus and I/O controller system. The NIC in the originating server is coupled to the NIC in the target server by the communication link, the HBA in the originating server is coupled to and controls the primary array, and the HBA in the target server is coupled to and controls the redundant array. Respective program processes reside in the servers’ respective system memories that control the servers’ respective CPUs so as cause the servers to carry out conventional synchronous data replication operations. The respective program processes residing in the respective system memories include respective operating system (“O/S”), RAID driver, replication driver, and network driver/communication stack processes. Application-level database processes also reside in the originating server’s system memory. (Specification, page 1, line 17 to page 2, line 3).

A human user of the originating server may issue a request, via a user interface to an application-level program process residing in the originating server, to store user data in a primary data volume in the primary array. In response to this request, the process generates a data write request to one or more O/S processes in the originating server that causes the originating server’s CPU to transmit the user data via the originating server’s system bus to, and

store the data in, the originating server's system memory. Thereafter, these O/S processes issue a request to originating server's RAID driver that causes the originating server's CPU to retrieve the data from the originating server's system via the originating server's system bus, to fragment the data into one or more data blocks (e.g., Small Computer Systems Interface (SCSI) protocol data blocks), to transmit the data blocks to, and store them in, the originating server's system memory, and to cause the originating server's I/O bus controller to retrieve the blocks from the originating server's system memory via the originating server's system bus and to forward the blocks to the originating server's HBA. The originating server's HBA then writes the data blocks in the primary data volume, and when the data blocks have been successfully written in the primary data volume, the originating server's HBA returns to the originating server's CPU via the originating server's I/O bus and controller system, and system bus a "write complete" message to indicate same. (Specification, page 2, lines 4 - 21).

Also in response to the request from application-level process, the originating server's O/S processes request that the originating server's replication driver processes command the originating server's CPU to generate a copy of the user data, and to fragment the copy of the user data into one or more replication data packets, and to store the packets in the originating server's system memory. After these packets have been stored in the originating server's system memory, the originating server's network driver processes cause the originating server's CPU to retrieve the packets from the originating server's system memory via the originating server's system bus, to encapsulate the packets with header and other information so as to form one or more Ethernet frames, to forward the frames to the originating server's NIC via the originating server's system bus, I/O bus and controller system, and to cause the originating server's NIC to forward the frames to the target server's NIC via the Ethernet link. (Specification, page 2, line 22 to page 3, line 2).

The target server's NIC receives the frames, and in response to the received frames, the target server's operating system, RAID driver, replication driver, and network driver/communication stack processes cause the target server's CPU to perform operations that de-encapsulate the data packets from the frames, reconstruct the user data from the data packets, fragment the reconstructed user data into one or more data blocks, and store the blocks in the redundant data volume. After all of the user data has been successfully written in the redundant volume, the target server's O/S, replication driver, and network driver processes cause the target server's CPU to perform operations that generate and forward to the originating server (i.e., via the target server's NIC and the communication link) a data replication success message frame indicating successful replication of the data in the redundant array. As can be appreciated, the above operations carried out by the target server's CPU involve a substantial number of transactions/interactions among, *inter alia*, the target server's CPU, system memory, and system bus. (Specification, page 3, lines 3 - 16).

In this conventional network, an excessively large number of CPU, system memory, and system bus transactions and operations must be performed in the originating and target servers to carry out the above data storage-related operations. Unfortunately, this causes excessive amounts of originating and target server CPU, system memory, and system bus processing resources to be consumed to carry out such operations. (Specification, page 3, lines 17 - 22).

Advantages of Embodiments of the Claimed Subject Matter

In contrast to the above conventional techniques, in embodiments of the claimed subject matter, the vast majority of operations that are performed in the originating and target servers to carry out data storage-related operations may be offloaded from the respective CPUs, system memories, and system buses in these servers to HBA I/O processors, internal HBA memories,

etc. This permits the number of CPU, system memory, and system bus transactions and operations that are performed in the originating and target servers to carry out data storage-related operations in these embodiments to be substantially reduced compared to the prior art. Advantageously, this prevents excessive amounts of originating and target server CPU, system memory, and system bus processing resources to be consumed when carrying out such operations, and frees such resources for use in other processing activities, and increases the speed and reliability with which such operations may be carried out in these embodiments, compared to the prior art. (Specification, page 23, lines 1 - 13).

VI. GROUND OF REJECTION TO BE REVIEWED ON APPEAL

In the Final Office Action, the Examiner raises two grounds of rejection. In the first ground of rejection, the Examiner rejected claims 1 - 32 and 39 - 43 under 35 USC § 102(b) as being anticipated by Ofek (U.S. 5,901,327). In the second ground of rejection, the Examiner rejected claims 33 - 38 under 35 USC § 103 as being unpatentable over Ofek. Appellant respectfully traverses each of these grounds of rejection, and appeals to the Honorable Board for review by the Honorable Board of each of these grounds of rejection.

VII. ARGUMENT

REJECTION UNDER 35 USC § 102(b) OVER OFEK

CLAIMS 1 - 32 AND 39 - 43

The Examiner has taken the position that claims 1- 32 and 39 - 43 are anticipated under 35 USC § 102(b) by Ofek. To anticipate a patent claim, a single prior art reference must disclose every limitation of the claim. *In re Schreiber*, 128 F.3d 1473, 1477, 44 USPQ2d 1429 (Fed. Cir. 1997).

As the Honorable Board is well aware, when determining the broadest reasonable interpretation to be given to the claims during *ex parte* examination in the U.S. Patent & Trademark Office, if no special definition has been clearly set forth in the Specification, the claim language is to be given its plain meaning **to those skilled in the art to which the invention most closely pertains**. *Toro Co. v. White Consol. Indus., Inc.*, 199 F.3d 1295, 1299, 53 USPQ2d 1065, 1067 (Fed. Cir. 1999); *Manual of Patent Examining Procedure* (MPEP) § 2111.01. That is, the words of the claims must be read as they would be interpreted by those skilled in technology of the subject application. In order to establish such plain meaning, it is well-established that the Honorable Board may resort to dictionaries and/or learned treatises of the pertinent technology. See id.

It is beyond reasonable dispute that the subject application and claims on appeal relate to the field of data storage technology. To those skilled in the art of data storage technology, the customary meaning of the language “host bus adapter” in claim 1 is an I/O adapter that connects a host I/O bus to the host’s memory system. Additionally, as used in the art, a “host I/O bus” is an I/O bus that is used to connect a host computer to one or more storage systems or storage devices. Moreover, an “I/O bus” means any path used for the transfer of data and control information between I/O adapters and storage controllers or storage devices. See, e.g., *The RAID Book*, Sixth Edition, Raid Advisory Board, Inc., 1997, pp. 257 and 258 (attached hereto in Appendix B: Evidence Appendix). Significantly, the Assignee of Ofek, EMC Corporation, actually is a member of the RAID Advisory Board (the publisher of *The RAID Book, supra*), fully supported its publication, and was a material contributor to *The RAID Book, supra*. See, *The RAID Book, supra*, copyright page (attached hereto in Appendix B: Evidence Appendix).

When the limitations of independent claims 1, 5, 10, 14, 16, 20, 25, 29, 31, 39, and 42 are construed in light of the above definitions, it is clear that Ofek cannot anticipate or render

obvious claims 1 - 32 and 39 - 43. For example, in the case of independent claim 1, Ofek nowhere discloses:

“A first adapter to use in a first server in a network, the first adapter comprising . . . a host bus adapter (HBA) including circuitry to cause, in response to a first request received by the HBA, execution of a first data storage-related operation associated with a first set of mass storage devices and to issue, also in response to the first request, a second request from the HBA to a second adapter in a second server in the network to cause the second adapter to perform, in response to the second request, a second data storage-related operation associated with a second set of mass storage devices. (Independent claim 1).

In contrast to these limitations of independent claim 1, in pertinent part, Ofek discloses a system in which a primary data storage system 14 is connected via a signal path 18 to a first host 12 at a first site, and a secondary data storage system 46 is connected to second host 52 at a remote site. The storage systems 14 and 46 comprise respective controllers 16 and 44. Controllers 16 and 44 comprise respective adapters 26 and 54 that are connected to hosts 12 and 52, respectively. See Ofek's Figure 1.

Additionally, the first adapter of claim 1 is used in a first server. Ofek nowhere discloses a first adapter of the type claimed in claim 1 that is used in a server. At page 2 of the Final Office Action, the Examiner asserts that the first adapter of claim 1 corresponds to controller 16 in Ofek's disclosed arrangement, and the first server of claim 1 corresponds to host A in Ofek's disclosed arrangement. However, contrary to the Examiner assertion, in Ofek's disclosed arrangement, controller 16 is not used in host A, but instead, as stated above, is comprised in data storage system 14. Also contrary to the Examiner's assertions at page 2 of the Final Office

Action, in light of, *inter alia*, the above definitions, controller 16 cannot be said to comprise a host bus adapter of the type described in claim 1. For example, controller 16 does not comprise an I/O adapter that connects a host I/O bus to the host's memory system.

These distinctions between claim 1 and Ofek's disclosed arrangement are more than merely academic. For example, the specific features of claim 1 permit embodiments of the invention of claim 1 to operate in a manner that is substantially different from Ofek's disclosed arrangement, and to achieve advantages that cannot be achieved by Ofek's disclosed arrangement. See, e.g., Specification, page 23, lines 1 to 13.

In pertinent part, independent claims 5, 10, 14, 16, 20, 25, 29, 31, 39, and 42 recite limitations of claim 1 and/or limitations that are similar, in substance, to the limitations of claim 1 that distinguish claim 1 over Ofek, and permit the invention of claim 1 to achieve these profound advantages over Ofek. See, claims 5, 10, 14, 16, 20, 25, 29, 31, 39, and 42. Thus, in view of the specific, advantageous combinations of limitations of the independent claims that are not disclosed or suggested in Ofek, it is respectfully submitted that Ofek does not anticipate or render obvious the claims of the subject application.

In the Advisory Action, the Examiner maintained the rejection of the subject application based on the following assertions:

The broadest reasonable interpretation of a "host bus adapter" includes those that are defined in well-established dictionaries and meanings that are well known by those of ordinary skill in the art. However, it does [sic, is] not limit [sic, limited] to Applicant's specific definition as defined by Applicant in the arguments rather than the specification. For example, Microsoft Computer Dictionary, Fifth Edition, defines host adapter as "a device for connecting a peripheral to the main computer, typically in the form of an expansion card. Also called: controller,

host bus adapter.” Ofek clearly discloses in FIG. 1 such a device (controller 16) for connecting a peripheral (storage device 20) to the main computer (host A). Therefore, the rejection of claims 1-44 [sic, 43] is deemed to be proper. Advisory Action, page 2.

In response, Appellant notes that the Examiner’s proffered definition of “host bus adapter” is taken from Microsoft Computer Dictionary, which is **not** a dictionary or learned treatise **in the field of the claimed invention (i.e., data storage)**. In contrast, Appellant’s definition of “host bus adapter” is from a dictionary comprised in a learned treatise in the field of data storage. Moreover, the Assignee of Ofek actually is a member of the publisher of Appellant’s learned treatise, fully supported its publication, and was a material contributor to the treatise! Clearly, under relevant law and U.S. Patent & Trademark Office policy, the Examiner’s proffered definition of “host bus adapter” is entitled to no evidentiary weight, while Appellant’s definition of that term is authoritative.

Additionally, in the Final Office Action, the Examiner asserts that the “first server” of claim 1 corresponds to host A in Ofek’s disclosed arrangement. However, based on the Examiner’s proffered definition of “host bus adapter” in the Advisory Action, the Examiner asserts that the host bus adapter of claim 1 corresponds to controller 16 in Ofek’s disclosed arrangement. Final Office Action, page 2 and Advisory Action, page 2. However, controller 16 is comprised in data storage system 14, not host A, in Ofek’s disclosed arrangement. See, Ofek’s Figure 1. **Thus, the Examiner’s position completely ignores the fact that independent claim 1 requires that the first adapter (1) be for use in the first server, and (2) comprise the host bus adapter.** See, independent claim 1. The other independent claims on appeal contain similar limitations. See, claims 5, 10, 14, 16, 20, 25, 29, 31, 33, 37, 39, and 42. Quite simply, in

proffering the Examiner's definition of "host bus adapter," and in construing Ofek and the claims on appeal, the Examiner has ignored limitations of the claims that cannot be reconciled with the Examiner's proffered definition and claim construction. This cannot serve as proper basis for rejecting the claims on appeal!

Thus, for the above reasons, among others, it is respectfully submitted that Ofek does not anticipate or render obvious claims 1 - 32 and 39 - 43. Thus, it is respectfully submitted that the Examiner's rejection of claims 1 - 32 and 39 - 43 under 35 USC § 102(b) as being anticipated by Ofek is in error, and should be reversed by the Honorable Board.

REJECTION UNDER 35 USC § 103 OVER OFEK

CLAIMS 33 - 38

As the Honorable Board is well aware, in order to establish a *prima facie* case of obviousness:

First, there must be some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings. Second, there must be a reasonable expectation of success. Finally, the prior art reference (or references when combined) must teach or suggest all the claim limitations. . . The teaching or suggestion to make the claimed combination and the reasonable expectation of success must both be found in the prior art, not in applicant's disclosure. *In re Vaech*, 947 F.2d 488, 20 USPQ2d 1438 (Fed. Cir. 1991). MPEP § 2143.

The teachings and deficiencies of Ofek vis-à-vis the claims on appeal are discussed in detail above. In contrast to Ofek, independent claim 33 recites:

Computer program instructions residing in a computer-readable memory, the computer program instructions comprising a set of instructions that when executed **by a first processor comprised in a host bus adapter (HBA) in a first server** cause:

execution of, in response to a first request, a first data storage-related operation associated with a first set of mass storage devices; and

generation of, also in response to the first request, a second request that may be issued from the HBA to a second processor in a second server to cause the second processor to perform, in response to the second request, a second data storage-related operation associated with a second set of mass storage devices.

(Independent Claim 33).

Thus, in claim 33, the first processor is comprised in a first host bus adapter in a first server. Ofek nowhere discloses a first processor that is comprised in an HBA in a server that is capable of causing execution of the operations recited in claim 33. These distinctions between claim 33 and Ofek's disclosed arrangement are more than merely academic. For example, the specific features of claim 33 permit embodiments of the invention of claim 33 to operate in a manner that is substantially different from Ofek's disclosed arrangement, and to achieve advantages that cannot be achieved by Ofek's disclosed arrangement. See, e.g., Specification, page 23, lines 1 to 13.

In pertinent part, independent claim 37 recites limitations of claim 33 and/or limitations that are similar, in substance, to the limitations of claim 33 that distinguish claim 33 over Ofek,

and permit the invention of claim 33 to achieve these profound advantages over Ofek. See, claims 37. Thus, in view of the specific, advantageous combinations of limitations of independent claims 33 and 37 that are not disclosed or suggested in Ofek, it is respectfully submitted that Ofek does not render obvious the claims 33 - 38 of the subject application. Thus, it is respectfully submitted that the Examiner's rejection of claims 33 - 38 under 35 USC § 103 as being unpatentable over Ofek is in error, and should be reversed by the Honorable Board.

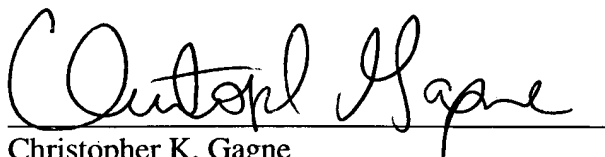
CONCLUSION

For the foregoing reasons, Appellant respectfully submits that the Examiner's final rejections of claims 1 - 43 are erroneous. Accordingly, Appellant respectfully requests that the Honorable Board of Patent Appeals and Interferences reverse the Examiner and direct that all of the currently pending claims be allowed.

Please charge any shortages and credit any overcharges to Deposit Account No. 50-0221.

Respectfully submitted,

Date: 20 Sept 2001



Christopher K. Gagne
Attorney For Appellant
Reg. No. 36,142
Telephone: 480-715-4055
Fax: 480-715-7738

c/o Blakely, Sokoloff, Taylor, & Zafman, LLP
12400 Wilshire Boulevard, Seventh Floor
Los Angeles, CA 90025-1026
(503) 684-6200

APPENDIX A: CLAIMS APPENDIX

1. A first adapter to use in a first server in a network, the first adapter comprising:
a host bus adapter (HBA) including circuitry to cause, in response to a first request received by the HBA, execution of a first data storage-related operation associated with a first set of mass storage devices and to issue, also in response to the first request, a second request from the HBA to a second adapter in a second server in the network to cause the second adapter to perform, in response to the second request, a second data storage-related operation associated with a second set of mass storage devices.
2. The first adapter of claim 1, further comprising:
additional circuitry to issue, in response to a first message from the second adapter, a second message to a process in the first server, the first message indicating that the second data storage-related operation has been completed, the second message indicating that a respective data storage-related operation requested by the first request has been completed.
3. The first adapter of claim 2, wherein:
the second request and the first message each comprise a respective target node address field, initiating node address field, command field, and message identification field; and
respective message identification fields in the second request and the first message contain identical respective values.
4. The first adapter of claim 1, wherein the circuitry and the second adapter each comprise a respective I/O processor.

5. A first adapter to use in a first server, the first adapter comprising:
- a host bus adapter (HBA) including circuitry to perform, in response to a request, a data storage-related operation associated with a first set of mass storage devices, the request being issued from a second adapter in a second server in response to another request received by the second adapter to cause the second adapter to perform, in response to the another request, another data storage-related operation associated with a second set of mass storage devices.
6. The first adapter of claim 5, further comprising:
- additional circuitry to issue a first message to the second adapter to indicate that the data storage-related operation associated with the first set of mass storage devices has been completed, and the second adapter is configured to issue, in response to the first message, a second message to a process in the second server.
7. The first adapter of claim 6, wherein:
- the request issued from the second adapter and the first message each comprise a respective target node address field, initiating node address field, command field, and message identification field; and
- respective message identification fields in the request issued from the second adapter and the first message contain identical respective values.
8. The first adapter of claim 1, wherein:
- the first adapter is coupled to the first set of mass storage devices;

the second adapter comprises a second host bus adapter coupled to the second set of mass storage devices; and

the first adapter and the second adapter are coupled together via a network communication link.

9. The first adapter of claim 5, wherein:

the first adapter is coupled to the first set of mass storage devices;

the second adapter comprises a second host bus adapter coupled to the second set of mass storage devices; and

the first adapter and the second adapter are coupled together via a network communication link.

10. A first input/output (I/O) processor, the first I/O processor being comprised in a host bus adapter (HBA) in a first server and being configured so as to able to execute a set of operations comprising:

execution, in response to a first request, of a first data storage-related operation associated with a first set of mass storage devices;

issuance, also in response to the first request, of a second request from the first server to a second I/O processor in a second server to cause the second I/O processor to perform, in response to the second request, a second data storage-related operation associated with a second set of mass storage devices.

11. The first I/O processor of claim 10, wherein:

the first server comprises the first set of mass storage devices;

the second server comprises the second set of mass storage devices; and
the first and second servers are coupled together via a network communication link.

12. The first I/O processor of claim 10, wherein:

the second data storage-related operation comprises one or more of the following operations:

designation of a first data volume in a second server in which data stored in a second data volume in the first server is to be replicated;

expansion of a size of the target data volume; and

replication in the first data volume of the data.

13. The first I/O processor of claim 10, wherein:

the second data storage-related operation comprises one or more of the following operations:

termination of a previously-established association between a first data volume in the second server and a second data volume in the first server, the association designating that data stored in the second data volume is to be replicated in the first data volume; and

re-establishment of the previously-established association after the previously-established association has been terminated.

14. A first input/output (I/O) processor, the first I/O processor being included in a host bus adapter (HBA) in a first server and being configured to execute a set of operations comprising:

execution, in response to a request, of a first data storage-related operation associated with a first set of mass storage devices; and

issuance, after completion of the execution of the first data storage-related operation, of a first message from the first server to a second I/O processor in a second server to cause the second I/O processor to issue, in response to the first message, a second message to a process in the second server, the second I/O processor being operatively configurable both to generate the request and to perform a second data storage-related operation associated with a second set of mass storage devices, in response to another request from the process.

15. The first I/O processor of claim 14, wherein:

the first set of mass storage devices comprises one or more respective mass storage devices;

the second set of mass storage devices comprises one or more respective mass storage devices; and

the first message is comprised in a frame.

16. A method of using a first adapter including a host bus adapter (HBA) in a first server in a network, the method comprising:

using the first adapter to perform, in response to a first request, a first data storage-related operation associated with a first set of mass storage devices; and

issuing from the first adapter, also in response to the first request, a second request to a second adapter in a second server in the network to cause the second I/O adapter to perform, in response to the second request, a second data storage-related operation associated with a second set of mass storage devices.

17. The method of claim 16, further comprising:

issuing from the first adapter, in response to a first message from the second adapter, a second message to a process in the first server, the first message indicating that the second data storage-related operation has been completed, the second message indicating that a respective data storage-related operation requested by the first request has been completed.

18. The method of claim 17, wherein:

the second request and the first message each comprise a respective target node address field, initiating node address field, command field, and message identification field; and

respective message identification fields in the second request and the first message contain identical respective values.

19. The method of claim 16, wherein the first adapter and the second adapter each comprise a respective I/O processor.

20. A method of using a first adapter including a host bus adapter (HBA) in a first server, the method comprising:

using the first adapter to perform, in response to a request, a data storage-related operation associated with a first set of mass storage devices, the request being issued from a second adapter in a second server in response to another request received by the second adapter, the second adapter being configured to perform, in response to the another request, another data storage-related operation associated with a second set of mass storage devices.

21. The method of claim 20, further comprising:

issuing from the first adapter to the second adapter a first message to indicate that the data storage-related operation associated with the first set of mass storage devices has been completed, the second adapter being configured to issue, in response to the first message, a second message to a process in the second server.

22. The method of claim 21, wherein:

the request issued from the second adapter and the first message each comprise a respective target node address field, initiating node address field, command field, and message identification field; and

respective message identification fields in the request issued from the second adapter and the first message contain identical respective values.

23. The method of claim 16, wherein:

the first adapter is coupled to the first set of mass storage devices;

the second adapter comprises a second host bus adapter coupled to the second set of mass storage devices; and

the first adapter and the second adapter are coupled together via a network communication link.

24. The method of claim 21, wherein:

the first adapter is coupled to the first set of mass storage devices;

the second adapter comprises a second host bus adapter coupled to the second set of mass storage devices; and

the first adapter and the second adapter are coupled together via a network communication link.

25. A method of using a first input/output (I/O) processor included in a host bus adapter (HBA) in a first server, the method comprising:

using the first I/O processor to execute, in response to a first request, a first data storage-related operation associated with a first set of mass storage devices; and

using the first I/O processor to generate, also in response to the first request, a second request that may be issued from the first server to a second I/O processor in a second server to cause the second I/O processor to perform, in response to the second request, a second data storage-related operation associated with a second set of mass storage devices.

26. The method of claim 25, wherein:

the first server comprises the first set of mass storage devices;

the second server comprises the second set of mass storage devices; and

the first and second servers are coupled together via a network communication link.

27. The method of claim 26, wherein:

the second data storage-related operation comprises one or more of the following operations:

designation of a first data volume in the second server in which data stored in a second data volume in the first server is to be replicated;

expansion of a size of the first data volume; and

replication in the first data volume of the data.

28. The method of claim 26, wherein:

the second data storage-related operation comprises one or more of the following operations:

termination of a previously-established association between a first data volume in the second server and a second data volume in the first server, the association designating that data stored in the second data volume is to be replicated in the first data volume; and

re-establishment of the previously-established association after the previously-established association has been terminated.

29. A method of using a first input/output (I/O) processor included in a host bus adapter (HBA) in a first server, the method comprising:

using the first I/O processor to execute, in response to a request, a first data storage-related operation associated with a first set of mass storage devices; and

using the first I/O processor to generate, after completion of the execution of the first data storage-related operation, a first message that may be issued from the first server to a second I/O processor in a second server to cause the second I/O processor to issue, in response to the first message, a second message to a process in the second server, the second I/O processor being operatively configurable both to generate the request and to perform a second data storage-related operation associated with a second set of mass storage devices, in response to another request from the process.

30. The method of claim 29, wherein:

the first set of mass storage devices comprises one or more respective mass storage devices;

the second set of mass storage devices comprises one or more respective mass storage devices; and

the first message is comprised in a frame.

31. A network comprising:

a first server associated with a first set of mass storage devices and including a first input/output (I/O) processor;

a second server remote from the first server, associated with a second set of mass storage devices, and including a second I/O processor;

a network communication link coupling the first server to the second server;

the first I/O processor being comprised in a host bus adapter (HBA) in the first server and being configured so as to be able to cause the following operations:

execution, in response to a first request, of a first data storage-related operation associated with the first set of mass storage devices;

issuance, also in response to the first request, of a second request from the first server to the second server via the link to cause the second I/O processor to perform, in response to the second request, a second data storage-related operation associated with the second set of mass storage devices.

32. The network of claim 31, wherein:

the second I/O processor is configured so as to be able to cause the following operations to be executed:

execution, in response to the second request, of the second data storage-related operation; and

issuance, after completion of the execution of the second data storage-related operation, of a first message from the second server to the first server via the link to cause the first I/O processor to issue a second message to a process in the first server to indicate a completion of the first data storage-related operation and the second data storage-related operation.

33. Computer program instructions residing in a computer-readable memory, the computer program instructions comprising a set of instructions that when executed by a first processor comprised in a host bus adapter (HBA) in a first server cause:

execution of, in response to a first request, a first data storage-related operation associated with a first set of mass storage devices; and

generation of, also in response to the first request, a second request that may be issued from the HBA to a second processor in a second server to cause the second processor to perform, in response to the second request, a second data storage-related operation associated with a second set of mass storage devices.

34. The computer program instructions of claim 33, wherein:

the first server comprises the first set of mass storage devices;

the second server comprises the second set of mass storage devices; and

the first and second servers are coupled together via a network communication link.

35. The computer program instructions of claim 34, wherein:

the second data storage-related operation comprises one or more of the following operations:

designation of a first data volume in the second server in which data stored in a second data volume in the first server is to be replicated;

expansion of a size of the first data volume; and

replication in the first data volume of the data.

36. The computer program instructions of claim 34, wherein:

the second data storage-related operation comprises one or more of the following operations:

termination of a previously-established association between a first data volume in the second server and a second data volume in the first server, the association designating that data stored in the second data volume is to be replicated in the first data volume; and

re-establishment of the previously-established association after the previously-established association has been terminated.

37. Computer program instructions residing in a computer-readable memory, the computer program instructions comprising a set of instructions that when executed by a first processor included in a host bus adapter (HBA) in a first server cause:

execution, in response to a request, of a first data storage-related operation associated with a first set of mass storage devices; and

generation, after completion of the execution of the first data storage-related operation, of a first message that may be issued from the first server to a second processor in a second server to cause the second processor to issue, in response to the first message, a second message to a

process in the second server, the second processor being operatively configurable both to generate the request and to perform a second data storage-related operation associated with a second set of mass storage devices in response to another request from the process.

38. The computer program instructions of claim 37, wherein:

the first set of mass storage devices comprises one or more respective mass storage devices;

the second set of mass storage devices comprises one or more respective mass storage devices; and

the first message is comprised in a frame.

39. A first server, comprising:

a host bus adapter (HBA) comprising a first processor;

the first processor being configured to be able to cause:

execution, in response to a first request, of a first data storage-related operation associated with a first set of storage devices, the first set of storage devices being associated with the first server; and

issuance, also in response to the first request, of a second request from the first server to a second server to cause a second processor in the second server to perform, in response to the second request, a second data storage-related operation associated with a second set of storage devices, the second set of storage devices being associated with the second server.

40. The first server of claim 39, wherein the set of storage devices comprises a set of one or more mass storage devices.

41. The first server of claim 39, wherein the second server is remote from the first server.

42. A first server, comprising:

a host bus adapter (HBA) including a first processor;

the first processor being configured to be able to cause the following operations to be executed:

execution, in response to a request, of a first data storage-related operation associated with a first set of storage devices; and

issuance, after completion of the execution of the first data storage-related operation, of a first message from the first server to a second processor in a second server to cause the second processor to issue, in response to the first message, a second message to a process in the second server, the second processor being operatively configurable both to generate the request and to perform a second data storage-related operation associated with a second set of storage devices, in response to another request from the process.

43. The first server of claim 42, wherein:

the first set of storage devices comprises one or more mass storage devices;

the second set of storage devices comprises one or more mass storage devices; and

the first message is comprised in a frame.

Appendix B: Evidence Appendix

Appellant relies upon the following in this Appeal:

The Title and copyright pages, as well as, pages 257 and 258 of *The RAID Book*, Sixth Edition, Raid Advisory Board, Inc., 1997. Copies of these pages of *The RAID Book*, *supra*, were entered into the record as Exhibit A to Appellant's Amendment After Final Rejection. A copy of said Exhibit A is attached hereto in this Evidence Appendix.

THE **RAID** book

A STORAGE SYSTEM TECHNOLOGY HANDBOOK

6TH
EDITION



<http://www.raid-advisory.com>



EXHIBIT A

by Paul Massiglia

Copyright © 1995, 1996, 1997 The RAID Advisory Board,
Inc.

The RAIDbook is published to further the understanding and use of RAID technology. Copies of the RAIDbook are available from The RAID Advisory Board, 13 Marie Lane, St. Peter, MN 56014, and from RAID Advisory Board member organizations. The RAID Advisory Board, Inc. hereby permits RAID Advisory Board member organizations to make and distribute copies of all or portions of The RAIDbook subject to the following restrictions:

1. All such copies must bear this copyright notice and legend.
2. Material from the RAIDbook must be reproduced without any alteration or modification.

The RAIDbook would not have been possible without the full support of the entire RAID Advisory Board, and in particular the following companies who contributed material for inclusion:

Array Technology Corporation

Chantal Systems, A Division of BusLogic, Inc.

EMC Corporation

Formation, Inc.

Hewlett-Packard

NCR Corporation

Quantum Corporation

Storage Computer Corporation

Digital Equipment Corporation deserves special recognition not only for material input to the RAIDbook, but also for editing and coordinating publication of the RAIDbook.

Finally, the Berkeley Papers (listed in Appendix 2) should be mentioned, as they represent the seminal work on RAID, and therefore influenced the RAIDbook significantly.

The RAIDbook was prepared using Microsoft Word Version 6.0a executing on the Windows NT operating system, Version 3.51 and Windows95.

CHRISTOPHER GAGNE

host-based array	A disk array whose <i>Array Management Function</i> executes in host computer(s) rather than in a disk system. The member disks of a host-based array may be part of different disk systems. <i>cf.</i> controller-based array .
host bus	A host I/O bus (q.v.)
host bus adapter	An I/O adapter (q.v.) that connects a host I/O bus (q.v.) to the host's memory system.
host computer	Any computer system to which disks are attached and accessible for data storage and I/O. Mainframes, servers, workstations and personal computers, as well as multiprocessors and computer complexes such as clusters and sysplexes are all referred to as host computers in RAID Advisory Board publications.
host I/O bus	An I/O bus (q.v.) used to connect a host computer to storage systems or storage devices.
hot disk	A disk whose capacity to execute I/O requests is saturated by an I/O load.
hot file	A frequently accessed file. Hot files are generally the root cause of hot disks (q.v.), although this is not always the case. A hot disk can also be caused by operating environment I/O, such as paging or swapping.
hot spare (disk)	A disk being used as a hot standby component (q.v.)
hot standby (component)	A redundant component in a failure tolerant storage system that has power applied and is ready to operate, but which does not perform its task as long as the primary component for which it is standing by is functioning properly. Hot standby components are used to increase storage system availability by allowing systems to continue to function in the presence of a failed component. When the term hot standby designates a disk, it specifically means a disk that is spinning and ready to be written to, for example, as part of a rebuilding (q.v.) operation.
hot swap	The substitution of a replacement unit (RU—q.v.) in a storage system for a defective unit, where the substitution can be performed while the system is performing its normal function. Hot swaps are manual operations typically performed by humans— <i>cf.</i> automatic swap, cold swap, warm swap .



- Independent access array** A disk array whose data mapping is such that different member disks can execute multiple application I/O requests concurrently. *cf.* **parallel access array**
- inherent cost** The cost of a system expressed in terms of the number and type of components it contains. The concept of inherent cost allows technology-based comparisons of disk system alternatives by expressing cost in terms of number of disks, ports, modules, fans, power supplies, cabinets, etc. Because it is inexpensively reproducible, software is generally assumed to have negligible inherent cost.
- I/O adapter** A hardware device that converts between the timing and protocol of a host's memory bus and that of an I/O bus. In the context of storage systems, I/O adapters may be contrasted with **embedded storage controllers** (q.v.) , which perform timing and protocol conversion functions as well as others such as device fan-out, cache, and RAID.
- I/O bus** Any path used for the transfer of data and control information between I/O adapters and storage controllers or storage devices. An I/O bus consists of cables, connectors, and all associated drivers, receivers, transducers, and other electronics required to make it function. *cf.* **channel**
- I/O bottleneck** Any resource in the I/O path (e.g., a device driver, an I/O adapter, an I/O bus, an intelligent controller, or a disk) whose performance limits the performance of a storage system or I/O system as a whole.
- I/O driver** A host computer software component (usually part of an operating system) whose function is to control the operation of peripheral controllers or adapters attached to the host computer. I/O drivers communicate between applications and I/O devices, and in some cases may participate in data transfer, although this is rare with disk drivers, since most disk adapters and controllers contain hardware that performs data transfers.
- I/O-intensive** A characterization of applications. An I/O-intensive application is one whose performance depends strongly on the performance of the I/O system that provides its I/O services.

Appendix C: Related Proceedings Appendix

As stated in Appellant's Brief, to the best of the knowledge of the Appellant, the undersigned attorney, and the Assignee of the subject application, there are no prior or pending appeals, interferences, or judicial proceedings which may be related to, directly affect or be directly affected by, or have a bearing on the Honorable Board's decision in the subject appeal. Accordingly, there are no documents to attach hereto as this Appendix.